Table of Contents

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 158705129X Publisher: Cisco Press
Print Publication Date: 2004/02/24

Chapter 2. Passive Optical Networks in the First Mile

As Internet traffic doubles every six months or so, there is a tremendous surge in end-user bandwidth requirements. 14.4-kbps modems were replaced by higher-speed modems, which in turn were replaced by digital subscriber line (DSL) and cable modems. However, even these advances could not alleviate the basic bottleneck in access networks. The bottleneck stems from two causes. First, at this time, there is ample capacity in the core and metro area due to the emergence of wavelength-division multiplexing (WDM) as a premier technology for high-speed transport; WMD maximizes the use of the near-infinite bandwidth offered by the optical fiber by sending multiple data streams on multiple wavelengths. Second, the ratio of sink users to sources in access is quite high. Moreover, home and work PCs are able to operate at gigabit-level speeds, often creating a void between the core networks and the PCs for seamless data flow. In other words, multiple end users are connected to a single terminating line, each trying to extract and squeeze every possible bit of data from the line.

These two effects create a bottleneck in the access area, which, of course, we refer to as the *first mile issue*. This term can be replaced with the *last mile issue* without change of meaning. Why the access area attributes such importance to revenue can be understood by the fact that it is the end users who generate the revenue on which the business chain of service providers, enterprises, and system and components vendors functions. To create excellent transport methodologies in the access area, newer technologies are tried and deployed. The business proposition of the access area can be understood by the fact that there are multiple end users, each generating small amounts of revenue. The amounts may be small, but the volume of end users is truly enormous, creating a very solid value proposition.

The first implementations of commercial enterprise solutions led to the deployment of broadband access, namely DSL, asymmetric DSL (ADSL), very-high-data-rate DSL (VDSL), ATM, and other solutions. However, each was limited by bandwidth and scalability issues among others. The advent of optical fiber as a means for transport of data at a low cost and high speed (bandwidth) led to showcasing fiber to end-user applications as a possible and pragmatic solution. Fiber to the enterprise or user deployment, although slow (due to the high initial deployment cost of fiber) is the best and possibly only way to circumvent the bandwidth bottleneck between the end user and the metro access network. If we consider the excellence in characteristics provided by an optical fiber in terms of its longevity and protocol transparency, we realize the long depreciation cycle that actually

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

justifies and in most cases lowers the cost of fiber as compared to legacy copper solutions. The solution to providing bandwidth to end users has to be low cost, efficient, easy to manage, and scalable (among other intricacies such as resiliency and interoperability and the solution must provide degrees of quality of service [QoS]).

Passive optical networks are a technological innovation that can alleviate the first mile bottleneck issue in access networks. As the name implies, passive optical networks are typically *passive*, in the sense that they do not have active components for data transport. They may be spread across different physical topologies. PON development, although propelled by the surge in bandwidth requirements, also answers a definite need for low-cost optical networks for end-user applications. During the initial phase of PON development, some of the primary desirable features of a PON were as follows:

- **Low-cost network, low-cost components**— Because the revenue was in the number of consumers (quantity) rather than the pure service delivery to each consumer (quality), the amount of investment each end user would have to make had to be kept to a bare minimum. In metro and core networks, at each network element a composite WDM signal could drop an entire wavelength or a group of wavelengths. In contrast, in the access first mile area, each network element at the consumer site had bit rates typically of the order of 100 Mbps or even less.
- **Ease of management**— The first feature of low cost also initiates the second point of management complexity. Management complexity creates undue surges in network equipment cost. What is desired is a simple, efficient, and scalable management system that can manage the network and guarantee the network users of some network parameters such as QoS, delay, fairness, throughput (service level agreement [SLA]), and resiliency.
- **Upgradability, in-service upgrade, and interoperability**— The rapid development of newer technologies creates a need for ease of upgradability in PONs. In soft upgrades, the basic fabric remains the same, but a software upgrade enhances the features of the network. Because most of the end users are residential customers or enterprise users, in-service upgrades are important so as not to disrupt real-time services in PONs. Finally, due to the high-volume nature of PON users, there is a strong probability of PON networks having multivendor equipment in them. To facilitate ease of communication and create fair-competition interoperability, standards must exist for PON network elements to talk to one another.
- **Guarantee of basic network features**— The PON must be able to guarantee the end users some degree of network parameters, which are promised at inception (SLA). Although low-cost networks and simple protocols are generally designed for best-effort service, the quantum leap from traditional broadband to PON represents a sea-change shift in the end-user paradigm, and end users must get the desired services through the PON to justify the cost of deployment.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
Print Publication Date: 2004/02/24

Network Profiling for Passive Optical Networks

In the preceding section we identified passive optical networking as the key to high-bandwidth delivery to the end user. The sudden growth (and later stagnation) of WDM networking led to a void in the access area. The access area needed high-speed bandwidth, at a low cost, and needed to be able to provide reliable connection through a robust management platform. Fiber seemed to be the choice of many for providing the high-speed links to the end users. However, the cost of optical-to-electronic interfaces at end-user premises was the inhibiting factor for the industry, slowing deployment of fiber-based networks in the access area. In the mid-1990s, universities in North America proposed a low-cost and flexible scheme for broadband access networking called *passive optical networks (PONs)*. The term PON generically means optical networks that are pervasive and passive in nature. The passivity arises from the design contrast with metro networks, in which switching, amplification, and regeneration of the optical signal is carried out en-route from a source to a destination. Therefore, PONs typically do not require active optical components, namely switches, amplifiers, and so on, which are the key to the upward-driven prices seen in generic optical networking.

Characteristically, a PON consists of low-cost components that are engineered to provide reliable high-speed communication. One of the main motivating factors for achieving low cost through PON is the network topology itself. Unlike in a metro core and a long-haul network, where rings and meshes are the dominant topology often leading to expensive network gear for routing of connections, in the access area the use of fiber is limited to tree formation only, thus creating a natural broadcast medium that does not need routing protocols per se (only for the tree).

The Network Topology

Physically, a PON is based on a tree topology. The main reason for the tree topology is the formation of the network whereby a single central office is communicating with multiple end users geographically dispersed yet clustered around in the same area. This means that the relationship (graphical) between the single central office and the multiple end users can be described as one-is-to-many or simply one-to-many. When we relate this kind of one-to-many semblance to optical networking, the topology apposite to our requirement is that of a tree, where there is a single root and multiple leaves connected through the branches. Translating this to access area networks, this means that there are primarily two kinds of network elements in PONs in tree formation: the first at the central office of the service provider and the second at each of the end users.

Figure 2-1 shows a generic diagram of a PON network. Multiple end users are connected to a single network element (NE) at the central office through a tree of optical fibers—hence the passive network is formed. The NE at each end user is called the optical network unit (ONU), whereas the NE at the central office is called the optical line terminal unit

(OLTU, or sometimes just OLT). The OLT is connected on one side to these ONUs and likewise is connected on the other side to the metro core network, invariably a Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) or Gigabit Ethernet (GigE) network that drops bandwidth of granularity as required in the passive network. The OLT has the strategic importance of interfacing between the first mile access components (namely the ONUs) and the metro core network. Note that the OLT is a rather complicated device having to sometimes do protocol transfer and management functions. The ONUs, on the other hand, are connected to end users through multiple methods. In one embodiment, ONUs may be interfaced to the end user through a user network interface (UNI). In another scenario, ONUs may be connected to homes using a wireless solution. Yet another implementation has ONUs connected to homes and small businesses using copper solutions and taking advantage of long-reach Ethernet (LRE) concepts.

Figure 2-1. PON Topology

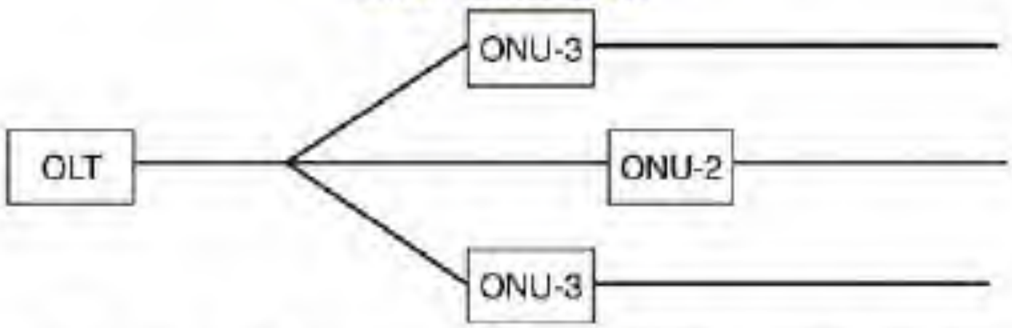

As mentioned previously, the structure of the network is that of a tree. The OLT is connected to the multiple ONUs through the fiber tree such that from the OLT a single fiber emanates toward the ONUs. The signal from the emanating fiber is split (in power) by an optical splitter. Multiple fibers emanate from the optical splitter, each carrying a portion of the signal from the ONU. Each of these multiple fibers is connected to an ONU. In some cases, these fibers may be further split, sharing two or more ONUs between them. Figure 2-2 explains this case of the tree topology. The manner of communication from the single OLT to the multiple ONUs is considered to be in the downstream manner, quite in conformance with the top-down approach cited in most networking environments. Also, the way communication is done, by using the single-to-multipoint splitter, which inherently is a passive device, sets about the motivation for the nomenclature of this network as a passive optical network. The reverse communication, that from the multiple ONUs to the single OLT, is considered to be in the upstream. Here again the power from multiple ONUs is added up in the splitter (which now acts as a combiner). Of course the story of this chapter is how to ensure upstream communication, the real problem in PON solutions, and ways to deal with it, because it can now be understood that successful upstream communication is the result of providing a collision-free environment when multiple ONUs transmit.

Figure 2-2. General Tree Topologies

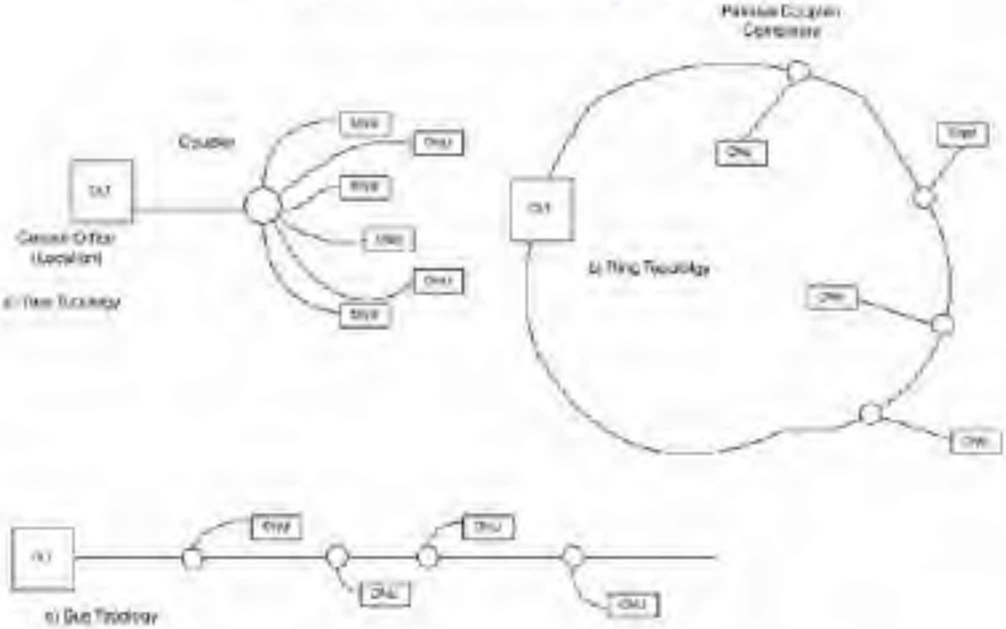

Passive Optical Networking: Problem Formulation

In the preceding section we covered the working of a PON and the main reasons for its nomenclature, the passivity of the medium. We now can digress to outline the main issue in PONs. We have noted and understood the importance of the star topology in passive optical networking and generally also for access networks: the need for connecting a large number of users to a single access point, namely the central office. This leads us to the obvious intuitive solution that a star topology is the best way to create passive optical networking. The passivity in the star-shaped network creates a bus kind of architecture, such that this bus is point to multipoint. This means if the central office (say connected to the center of the star) sends a signal into the star, all the end users get a copy of the signal. This is because the splitter splits the optical signal in such a way that each end user gets a replica of the signal.

Optically speaking, this is done as follows: Assume the splitter has a ratio of 1:N such that there is an input fiber and N output fibers connected to N users, and $N > 1$. The splitter is formed by the hot-fusing of the N output fibers with the input fiber such that when an optical pulse with some power P arrives, the power is almost evenly split into the N output ports, so that each of the N output ports gets P/N amount of power. Of course the power level drops down, but each port still gets a replica of the same signal. We use the following equation to calculate the optical power at any of the output ports of an optical splitter. If

P is the input optical power, and there are N output ports (hence N users), further if Δ is loss of signal in the coupler, the power at any of the N output ports is as follows:

$$p_{port} = (p - \triangle) - 10\log_{10} N$$

The authors developed the preceding equation, and the reader is encouraged to prove its correctness through analytical methods.

Earlier in this chapter, the behavior of a PON network was described as in a downstream direction—that is, from the central office to the end user. We also make special mention of the fact that the direction is downstream and, most importantly, the passivity creates a broadcast medium, one that is specially adapted for random traffic having broadcast nature, such as Ethernet.

In contrast, when we consider the other direction of communication—namely, from the end user to the central office (ONU to OLT in the previously defined nomenclature)—we see that the system is slightly more complicated: We observe that there are now N users (N ONUs) connected through our passive splitter (now acting as a combiner) to the central office (OLT). We see that the medium is to be shared in a ubiquitous way between the N users.

In contrast to the downstream case where the broadcast nature made communication very simple, the upstream case has the biggest problem, that of avoiding possibility for collision. Assume that any two of the N users want to send some data to the OLT. Now these two users do not know whether the other wants to send data in the time period they individually want to send. In other words, if users A and B want to send data to the OLT in a time period starting from t, both A and B do not know whether the other node has data to send commencing from t. This creates a possibility of collision between the transmissions from A and B in the same time slot starting from t. This means that in a PON network, there is an absolute uncertainty of transmission in the upstream direction; one uncertainty stems from the passive nature, and one needs a degree of intelligence to be avoided. A point to note here is that, unlike in wireless communication, detection of collisions in fiber is not feasible. Experiments have shown that although it may be possible to detect collisions by comparing the power of the collided signal with the initial signal power, it still is not a foolproof method for collision detection. This showcases the need for an efficient method to alleviate the problem of upstream communication, that of being able to transport data from ONUs to the OLT in a collision-free environment.

The previous discussion raises an important point about PONs: the need for a protocol for upstream communication—a protocol that guarantees collision-free upstream communication yet also ensures that each node gets a fair share of the bandwidth. There are numerous methods of providing upstream communication in the PON. WDM, time-division multiplexing (TDM), statistical TDM (STDM), and hybrid solutions are some of

First Mile Access Networks and Enabling Technologies By Adnan Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

the approaches. The most logical (but not cost-effective) solution to be considered for upstream communication is to use separate noninterfering channels for communication from the ONUs to the OLT. The obvious way to implement such a system is to have each ONU communicate to the OLT on a separate optical wavelength channel. This means each ONU would have its own operating wavelength, in which a laser would be tuned—one that emits modulated data in optical format. This also means that at the OLT side, for receiving data from N ONUs we would need N receivers (photodiodes) coupled with filters that filter just the desired frequency. Although this whole system is very effective, there is minimal deployment as of today and as projected for the near future. The reason for its minimal deployment is simple: cost. Using WDM (the method of multiplexing multiple wavelengths in the fiber), the system requires an arbitrarily large number of optical components that are generally very expensive. For example, an N-node network having N ONUs connected in star formation to N OLTs, needs N optical receivers and filters at the OLT, in addition to N optical transmitters, each at a different wavelength, one each at every ONU. This makes the system cost exorbitantly high, and there is particularly no justification in access networks, especially the first mile area, for such high-cost solutions.

The other two generic solutions proposed are using TDM and STDM. First, let's consider TDM. To multiplex more and more voice signals in a trunk, TDM was invented, whereby using high-speed sampling, one could multiplex multiple signals in the time domain by spatially separating the sampled signals. In other words, N number of signals could be delivered through just one channel by slotting the channel in N recurring slots and writing the data of the first signal in the first slot and so on, thus creating a time-division system. One of the benefits of such a system is the requirement of comparatively fewer components.

Let us now analyze the requirements of a TDM system for Ethernet over PON (EPON) upstream communication. Assume N end users (ONUs) are connected to an OLT through a passive star. The upstream channel is a single optical wavelength able to support the cumulative bit rates of the N nodes. The upstream channel is slotted such that one slot is destined for one node, meaning that if there are N nodes, node 1 gets slot 1 and the next slot it receives is slot N + 1 and so on. In other words, the periodicity of slots is N. When the system is functional, the OLT transmits data on these slots, which are predefined. A node detects a slot only if the slot is meant for the node. Otherwise the node just discards the slot. Therefore, it's imperative that the node knows when a slot starts and when it ends. This means the NEs (the N ONUs and the one OLT) need to be synchronized. Synchronization in PON networks for TDM is essential for successful communication. Synchronicity can be provided by using a standard clock reference (at the OLT, for instance); all the ONUs align their clocks accordingly. Note that while aligning to the OLT, the ONUs may be at different distances, and hence there may be variable delay from the OLT to the ONU, creating a slight difference in synchronization. This can be rectified by using mathematical tractable solutions to the far-near problem—a case when two users

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: [illegible]

communicate to an OLT, but both are at different distances. Hence, the closer user gets a packet ahead in time to the farther user, which creates a discrepancy.

Moreover, the very reason for deployment of PONs worldwide was to facilitate the huge surge in data networking, propelled by the absolute exponential growth of IP-centric traffic. It has been seen that IP traffic is bursty in nature, meaning that the interarrival times of IP packets is not uniform, and further that when IP packets arrive they come in large numbers followed by large time gaps, in which fewer or no packets at all arrive. This kind of traffic behavior means that bandwidth requirements of end-user applications are dynamically variable. This leads to the requirement of packet-oriented communication that supports IP traffic and one that can be made to be flexible to support variable traffic needs. Therefore, if we were to have a TDM system with fixed-sized time slots, one each for every node, this would not be efficient, because of the bursty behavior of IP traffic creating many voids. In other words, a large number of slots are empty because there is no data to send, and for nodes that have a high volume of data (burst), the periodic nature of the slots does not allow the nodes to cater to IP-centric traffic. This means that a TDM formation—that is, allotting slots of a TDM system—is not effective for a PON solution. Another issue is the cost involved in synchronizing the nodes of the PON. This is the motivation for a slight variant of TDM called STDM.

Consider a case in which six ONUs are connected to an OLT and are named A, B, C, D, E, and F. Now assume that a burst of Ethernet frames at the OLT of time duration t_1 seconds arrives and is to be routed to node A. In a TDM scheme, the burst would be broken into slots proportional to the TDM slot size and scheduled periodically. Of course, we note that in a passive broadcast architecture as that of the PON, when the OLT sends any information all the nodes receive the information (optical signal) and only the node(s), for whom it is meant creates an address matching and decodes the information. Coming back to the TDM system, it would take $t_1/T + Kt_1/T$ seconds to transmit the burst of information t_1 (burst size), where T is the TDM slot size, and K is the periodicity (after how many seconds does the slot for the node come back). This means that the system is inefficient on account of the long wait it would have and would typically need large buffers to fill in the waiting time at the OLT. Consider if t_1 is 10 ms, slot size T is 1 ms, and the periodicity is 3 ms, the total time needed to transfer the data out of the OLT is 40 ms. Now further assume that the system is such that there is no data for other ONUs at the OLT; this means that most of the downstream bandwidth is not occupied. To alleviate this efficiency issue in PON communication, two logical proposals led to communication protocols: Transmission Upon Reception (TUR) and Interleaved Polling with Adaptive Cycle Time (IPACT) (Gumaste and Chlamtac). The protocols are discussed in the concluding sections of this chapter, but before that the generic solution—the STDM approach—needs some attention. The bandwidth in the PON—assume for downstream communication—was slotted in periodic fixed time slots in the TDM approach. Unlike this scheme, in the STDM approach the downstream bandwidth is also slotted but the slots are not of fixed duration and, hence,

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24
Prepared for Fernando Silva, Safari ID: [illegible]
User number: [illegible]

are not periodic. The bandwidth is slotted into asynchronous variable time slots, such that these slots can be made to fit the size of a burst of frames/packets to cater to the IP-centric nature of data traffic. The immediate performance gain by using STDM is the ability to cater to bursty or any other traffic distribution and to be able to provide the much-needed dynamic flexibility for bandwidth-killer applications such as bandwidth on demand, video distribution, and so on. It is easy to understand both logically and intuitively the superlative performance benefits of using an STDM kind of scheme rather than a generic TDM scheme, but we have to also note the level of difficulty in implementing the scheme. Before we talk about the negative aspect of an STDM scheme, if any, however, we should look at its positive aspect and why it is a most natural contender for PON application.

STDM represents a seamless approach for first mile networks because of its flexible nature in being able to allocate bandwidth on demand, especially considering the randomness of the requirements for a large fraternity of consumers.

By removing the constraints of fixed-size time slots, in an STDM system we also remove the constraints associated with the need for node synchronization. This can prove particularly important from a cost perspective, especially as data rates increase and there is a need to use synchronization (leading to high-speed and high-performance clocks, at each and every node in the network).

Finally, an STDM solution is particularly necessary for honoring SLAs between users and service providers; in an STDM solution, the bandwidth allocation can all be either dynamic or pre-allocated, depending on the SLA between the user and the provider.

As this discussion makes obvious, an STDM solution is far better than a TDM solution in terms of cost, performance, and appropriateness for IP-centric traffic. However, the biggest issue is the deployment technique of an STDM solution.

Consider Figure 2-3, in which upstream communication and variable-length bursts of data are being sent to different nodes from the OLT. The issue here is how the ONUs (nodes) can know the start and end of the bursts, and, second, how do we allocate this bandwidth dynamically. Finally, we also need to outline a method for upstream communication, the initial problem seen in basic PON deployment.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 1587051298 Publisher: Cisco Press
Print Publication Date: 2004/02/24

Figure 2-3. Exemplifying the Need for Synchronization

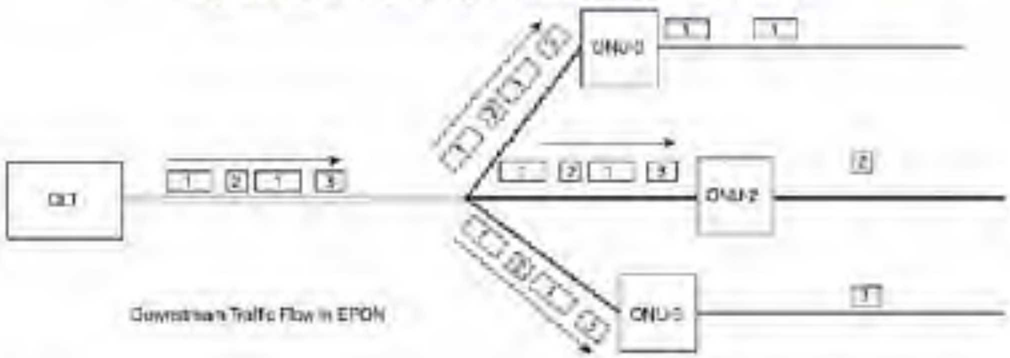

These issues are resolved later in this chapter when we cover the implementation of TUR (a protocol for EPON communication). Before digressing further, however, it's important to identify, define, classify, and compare the various types of PONs.

PON Classification—APON and EPON

As of today, there are three documented types of PONs: APON/BPON, GPON, and EPON. APON means ATM over PON, BPON stands for broadband PON, GPON means gigabit PON, and EPON is called Ethernet over PON. APON and BPON are variants of each other, whereas GPON is a vendor-proposed implementation not standardized as of this writing (but it is expected to be so soon). EPON, on the other hand, is the most significant type of PON that we discuss in detail as we progress through this chapter.

APON, the voice-centric PON prominently generating revenue in networks worldwide, evolved into a network that facilitated voice-type communication. In this regard, ATM over PON was a promising approach due to its circuit-centric nature. Through a consortium of carriers and vendors, the International Telecommunication Union (ITU) issued a series of recommendations—namely G.983.1, G.983.2, and G. 983.3—for possible deployment of APON (BPON) solutions. The recommendation of the BPON network was fully discussed in 1996 with regard to the formation of Full Services Access Network (FSAN). (You can find more information about this at http://www.fsanet.net.)

NTT and BellSouth (two carriers) formed the first collaborative effort for the deployment of BPON in 1998. This effort was called the *Common Technical Specification* (CTS). The number of partners finally joining hands grew to five when NTT and BellSouth were joined by British Telecom, France Telecom, and SouthWestern Bell Company (SBC) (each of these three being service providers).

BPON essentially means the same thing as APON. ATM is used because it can guarantee a good degree of QoS to the end user. The ATM layer resides on SONET interfaces, which are fed to an OLT that is connected through a tree to multiple ONUs. In other words, it is ATM over SONET in the last mile. ATM guarantees a committed bit rate, and SONET gives

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
ISBN: 1587051298 Publisher: Cisco Press
Print Publication Date: 2004/02/24

the necessary reliability and resilience for high QoS applications. The data rates seen in BPON are SONET/SDH equivalents—namely, 155 Mbps and 622 Mbps (OC-3 and OC-12, where OC means optical circuit).

In the upstream direction, the basic frame contains 53 cell slots of ATM, each 56 bytes long. Each cell slot consists of a 53-byte ATM cell and a 3-byte overhead. The overhead consists of a guard time, a preamble, and a delimiter. In the downstream direction, the basic frame consists of 56 cell slots, and each is 53 bytes long (same ATM format). In the downstream direction, each ONU gets the full complement of ATM cells but discards those cells (after making an address match) that are not destined for the incumbent ONU. In the upstream direction, the OLT gives ONU permission to send data by using a "grant" via the downstream cell. This cell is the physical layer operation, administration, and maintenance (PLOAM) cell. Despite the variance in the distances between different ONUs from the OLT, we have to ensure a mechanism that creates a collision-free communication. We outlined this as the far-near problem. This is solved in the following way: The OLT measures the distance to each ONU and then instructs every ONU to insert an appropriate delay so that all OLT-ONU distances are now virtually the same. This act of balancing the distances and solving the far-near problem is called *ranging*.

Advantages of BPON

The BPON architecture and paraphernalia represented the first way to implement passive optical networking for end users. What should be noted here are the motivating factors for BPON: First, there was a need to lower costs; second, the system needed to exhibit absolute protocol friendliness. On one hand, we see the absolute abundance of SONET/SDH networking and the capability to guarantee high-performance data (and voice and video) transport over a SONET network using its packet adaptation—namely, Packet over SONET (POS) (creating an amalgamation of ATM and SONET) in the core of networks. At the same time, and encouraging the initial deployment of ATM, there was a surge in the use of DSL and similar technology in the first mile, requiring a common denominator (ATM) that could cater to voice traffic as well as to bursty data traffic (of course, to a certain extent, because of its inherent TDM nature). These two factors prompted the need for APON and hence gave rise to the concept of BPON.

Note that despite the backing of multiple vendors and carriers in the initial deployment of BPON, several drawbacks have forced BPON to be deployed only partially; therefore, BPON cannot be seen as the best possible PON solution.

Of the many BPON drawbacks, the most relevant is the cost. BPON deploys primarily electronic-based technology—namely, SONET and ATM. Although effective and well proven, the problems with ATM and SONET technology are the associated cost and the heavy reliance on being a total electronic solution. The former creates a situation in which the economics do not support or validate BPON. The latter—the use of high-cost, low-

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

performance electronic systems—contrasts negatively with the advantages of passive optics. The overdependence on precise high-speed electronics means that the cost of deploying a BPON network is high. As indicated in Chapter 1, " Introduction to First Mile Access Technologies," the solution proposed for first/last mile access must have a low cost, one that suits the many and varied consumers who make up the revenue pool. The attempt to keep costs low is to a great extent nullified by the use of an electronic solution such as ATM and SONET. The higher costs significantly shrink the market for BPON from all end users (residential and small enterprises) to just small enterprises and homes with very high-capacity requirements (which are minimal). Therefore, despite its capability to technically solve the first/last mile problem, BPON is limited by cost. The other issue, that of not being able to use the passive optics, is another serious drawback for the deployment of BPON. PONs in a star topology represent a very viable solution for first mile networking. High costs associated with optical networks per se severely limit their use in the core and backbones. However, PONs represent one way to bring this cost down, because they use smarter technology and take advantage of the optical characteristics of coupler-like components that are technologically mature and extremely inexpensive. This direct advantage is nullified by the electronic bias of a BPON solution, which results in an expensive and inflexible system.

Gigabit PON

In 2001, the FSAN group (http://www.fsanet.net) initiated a new effort to standardize a high-bit-rate PON, now known as GPON. Apart from the need to support higher bit rates, the overall protocol was opened for reconsideration. The solution needed to be the most optimal and efficient in terms of support for multiple QoS-associated services such as voice and video. In other words, the GPON standardization process is basically a carry forward from the circuit-based processes seen in BPON, with some degrees of correction as compared to the earlier BPON version. Note that GPON uses Generic Framing Procedure (GFP) for framing and, hence, transport over the optical medium. The advantage of GFP is that multiple services can be mapped and adapted into a single (and generic) frame. This leads to less processing than other formats (such as ATM) and also enhancement of bit rates through forward error correction (a coding method to reduce transmission losses).

As a result of this FSAN activity, a new solution (that is, GPON) was showcased in the optical access area offering higher bandwidth (in the gigabit range) and enabling the same transport services (namely, circuit-based services such as voice).

As part of the GPON effort, a gigabit service requirement (GSR) document is implemented based upon the collective requirements of all member service providers (representing the leading Regional Bell Operating Companies [RBOCs] and Incumbent Local Exchange Carriers [ILECs]). The document is currently under submission to the ITU under the title "G.GPON.GSR."

First Mile Access Networks and Enabling Technologies By [illegible]
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: [illegible]

The salient features of GPON, summarized by the groups proposing GPON, are as follows:

- Full-service support, including voice and variants (TDM SONET and SDH), for data services through 10/100BASE-T
- Physical reach of at least 20 km, with a logical reach able to support 60 km (with regeneration)
- Support for various bit-rate options using the same protocol, including symmetrical and asymmetrical implementations using combinations of 622 Mbps, 1.25 Gbps, 2.5 Gbps, and so on
- Strong operation, administration maintenance, and provisioning (OAM and P) capabilities
- Security at the protocol level for downstream traffic (because of the multicast nature of PON)

When the vendors supporting such a hierarchy originally proposed GPON, multiple promises motivated its development:

- Higher data rates than previous versions of PON
- Better efficiency (typically for protocol considerations)
- GFP encapsulation of any type of service on standard SONET frames (125 ms)
- High efficiency with no overhead transport to native TDM traffic
- Dynamic allocation of upstream bandwidth via bandwidth maps (pointers) for each ONU

Let's digress from the GPON issue to consider GFP, specifically the importance attributed to GFP for optical transport networks. GFP provides a generic mechanism to adapt traffic from higher-layer signals over a transport network. The transport network could be of any type (such as SONET). Client signals may be packet based (IP) or circuit based. GFP has been officially standardized in the ITU document "G.7041." Because GFP provides a generic mechanism to facilitate the transport of multiple services in an efficient manner, over SONET, it ideally suits the basic GPON system, which is also an optical network with nodes synchronized with each other.

Briefly, GPON design characteristics are as follows:

- GPON is a frame-based, multiservice transport over PON.
- Upstream bandwidth allocation is accomplished via slot assignment.
- Database reports, security, and other issues are integrated into the physical (PHY) optical layer.
- Line coding is of a nonreturn to zero (NRZ) type, and therefore there is no need to return to zero after every pulse (1/0, for instance).
- GPON supports asymmetric line rate (namely, up to 1.25 Gbps).
- The upstream burst mode preamble includes clock and data recovery (CDR).

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

- The PHY layer also supports QoS.

Emergence of EPON—An Effective PON Solution, Particularly for IP-Centric Communication

The advent of data networking led to the emergence of a protocol that allowed computers to be linked and to "speak" using a universal and simple method. This method, Ethernet, was gradually incorporated into 95 percent of networks worldwide (initially by LANs but now by core networks, too).

SONET defines a standard for the transmission and multiplexing of data in the time domain through a network layer. SONET/SDH was created to cater to large volumes of voice traffic over circuit-switched networks and was standardized in the 1980s. SONET/SDH uses TDM and provides users a guaranteed level of service by assigning virtual channels within each frame. This guaranteed service ensures users that they can meet the critical requirements of voice traffic (namely, QoS), creating a platform for excellent timed delivery. Over the years, SONET has evolved to carry both voice and data traffic, and SONET data rates have increased from the original EC-1 (51 Mbps) to OC-192 (10 Gbps) today. Lately, some variations of SONET/SDH for data networking have been proposed (such as Packet over SONET [POS]), typically to match the growth from Ethernet.

It was very natural for vendors to consider SONET and SONET-like circuit-switched technologies (ATM, for example) as possible candidates for PON. This solution of circuits over PONs is good but is not the best. Why?

EPON has a relationship analogous to fish and water. First, the passivity in PON is the most conducive technical environment for Ethernet. Second, the cost involved in SONET networking was extremely high. The expensive electronics at each node and the level of precision required made SONET networking a very difficult proposition for first mile problems. In contrast, Ethernet was a ready solution; it was low cost and had good performance, especially for IP-centric communication (for bursty traffic arrivals). Ethernet also is a ubiquitous and well-understood topology, and it exhibits a plug-and-play nature that makes management and installation easier.

Despite these direct advantages, some industry pundits regard Ethernet in the access (and in the core) as a disruptive technology, viewing it with suspicion and deploying it only when alternatives dry up. We discuss in the following section the disruptive aspect of Ethernet and why it can withstand as a superior technology despite claims by legacy networking companies and pundits.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 9781587050862 Publisher: Cisco Press
Print Publication Date: 2004/02/24

Ethernet as a Disruptive Technology

Legacy service providers that have generated income through SONET/SDH circuit-switched technologies (primarily), which entail synchronization and expensive management—often criticize the deployment and development of packet technology, specifically because it is not revenue yielding.

Casson, Christensen, and others published a seminal paper titled "Ethernet in the MAN-A Case Study in Disruptive Technology" explaining the idealistic nature of SONET/SDH and ATM and how Ethernet proves to be a technology that is an able replacement for circuit-switched networks, especially in the first mile.

Performance oversupply is a principle that holds that mainstream technologies in certain industries improve at rates faster than the market can absorb. In Casson's paper, this is illustrated by the fact that this oversupply of new technology causes a paradigm shift in the market, creating a situation in which customers are given a wide choice of products. Christensen states that performance oversupply is the key prelude to signaling a future change in the basis of competition. Recent circuit-switched network performances have focused on increasing the distance and speed of WANs. As a result, these advances have overshot the needs of the access area. Figure 2-4 shows the pricing advantage of Ethernet as compared to a SONET (TDM) hierarchy.

Figure 2-4. Pricing of Ethernet and SONET

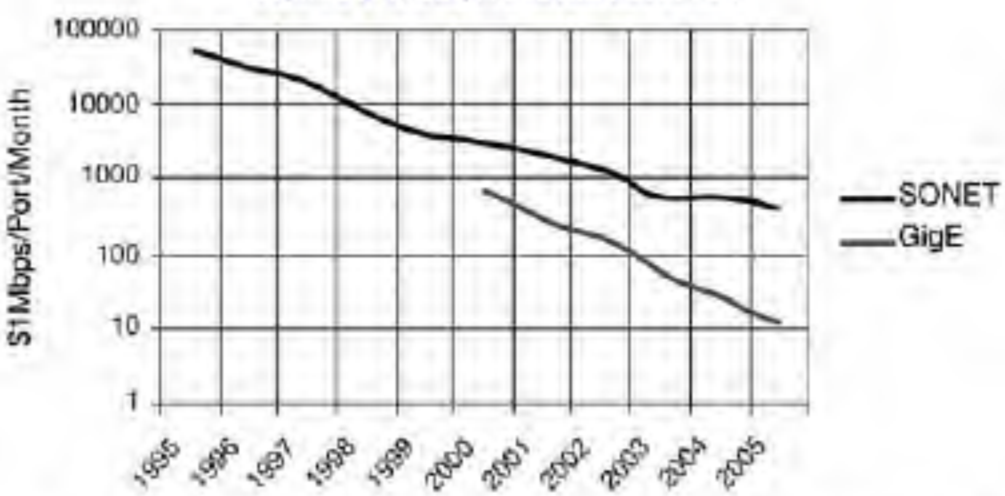

Most incumbent firms fail to realize the benefits of an emerging, disruptive technology. The ones that do are generally smaller startup kind of companies that do not focus on mainstream markets but create their own area of dominance (away from Christensen's theory of "resource dependence"). They instead find niche markets in which the attributes of the new technology find business value. Invariably, these attributes are the same ones that make it less attractive to mainstream customers. Consider, for instance, the LAN

market. Ethernet was first introduced into the LAN in the early 1980s by a number of small startups that initially focused on the DEC minicomputer market, and later the PC market. However, the large incumbent firm, IBM, delayed introducing its Token Ring technology to the LAN market until 1986. By that time, Ethernet already had an installed base of 30,000 networks containing at least 417,000 nodes. In the case of MANs and access area, Ethernet is also emerging in niche market segments where its attributes are valued. (Consider the metro Ethernet effort in Japan, for instance, an area technologically very superior to most parts of North America.) One MAN niche targeted by these companies has been small and medium-sized businesses and access area. Unlike large enterprises, access area generally does not require circuit-switched high-quality performance and reliability and values cost and flexibility in provisioning over other attributes. Another market niche targeted has been the largest metropolitan areas in the U.S. These large cities have a high density of potential customers and a large amount of unused, installed fiber that can be leased at a relatively low cost. The third and probably most important attribute is the management aspect and ability to provide service on demand. We all know the nuances of DSL and how difficult it is to provision a new service over an existing one without several software and hardware plug-ins. Circuit-switched technology inherently causes a lot of management problems, especially in provisioning networks. Packet-oriented networks, such as Ethernet, based over good topologies, such as PON, facilitate high-speed provisioning in the network domain.

Finally, disruptive technologies usually have lower profit margins than mainstream technologies. This factor often makes them unattractive to incumbents, whose cost structure and market valuation are based on higher-margin products. The low margins and profits, and associated lack of incumbents in the market, are the key factors that make the ensuing disorder and dislocation so pronounced.

All this leads to an important conclusion: Technologically, Ethernet is a superior candidate as compared to its circuit-switched counterparts. Financially, Ethernet costs much less to deploy than circuit solutions. This means that the end consumer (the end user) benefit immensely from the deployment and promulgation of Ethernet as a method for providing broadband access in terms of both cost and performance. It is the carriers and providers who need to refocus and relay cost-effective and performance-oriented strategies in the last mile. At times service providers may be dissatisfied with the costing structure (perhaps because of a temporary loss of revenue). However, it's in the best interest of the incumbent to incorporate the change and be able to sustain business for a long time than not to incorporate the change and face being wiped out of business altogether. It is in this context that EPON can be seen as the premier enabling technology, creating a very market-conducive technologically superior solution to solve the first mile access problem.

First Mile Access Networks and Enabling Technologies By Abhijit Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24
Prepared for Fernando Silva, Safari ID: [illegible]
User number: [illegible]

Standardization Efforts in EPON—The EFM Push

Because EPON appeared to be a natural and performance-oriented solution for the first mile access, a need arose to standardize the networking environment, especially the multiple subsystems that play a part in the deployment of EPON. In 2002, a task force called Ethernet in the First Mile (EFM) attempted to standardize EPON. The efforts of EFM culminated in a new standard for point-to-point as well as point-to-multipoint communications in the first mile using Ethernet as the central mechanism. Supported by the IEEE and to be nomenclated as IEEE 802.3ah, the EFM group's standardization has three main objectives:

- Point-to-point (P2P) Ethernet over fiber
- P2P Ethernet over copper for the last mile
- Point-to-multipoint (P2MP) Ethernet over fiber (as in PON)

The preceding three objectives are very important in the last mile and are discussed in detail in the following sections. The second objective—providing Ethernet over copper—is sometimes also considered the long-reach Ethernet (LRE) issue, which attempts to address the following considerations:

- Provide a family of physical layer specifications for the following:

    - 1000BASE–X (for transmission up to 10 km over single-mode fiber [using GigE])

    - 1000BASE–X (providing temperature extension in optics)
- Support far-end OAM in subscriber access networks (making sure there is excellent delivery of data to the end user)

    - Remote failure indication (ONU failing)

    - Remote loopback (MAC acts like a switch, looping back on the client side)

    - Link monitoring

The basic P2P architecture can be considered for seamless delivery of high-speed data over fiber. The application can be enterprise-level connectivity or transporting from an OLT to a central splitter.

Among others, some of the strategically important issues that have been upheld by EFM are as follows:

- **Finalization of upstream and downstream wavelengths**— In the upstream, a lower wavelength (for instance, 1310 nm or 1490 nm) is generally chosen. In the

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
Publisher: Cisco Press
Print Publication Date: 2004/02/24

downstream, however, a 1550-nm wavelength is chosen. The reason being that, first, it is better to have lower-cost equipment at consumer premises, which means 1310-nm lasers rather than the 1550-nm ones (and therefore a cost savings). Second, by choosing different wavelengths for upstream and downstream communication, we best utilize the same fiber without optical signal collision.

- **Temperature aspect**— Most ONUs, and for that matter any first mile customer apparatus, are kept at sites without indoor temperature controls and such. Therefore, the ONUs (and other apparatus, especially for optical components) must be able to operate in a wide range of temperatures without significant performance change. Laser instability because of temperature variation is a chief drawback to optics. An associated chirp (wavelength variation) of a laser's emitted light is dependent on temperature stability. Newer systems in the first mile have to take into account this chirp and develop a way to avoid rapid and large changes in frequency that may result from temperature changes.

Multipoint Application

Figure 2-5 shows the extension to P2MP networks.

Figure 2-5. Multipoint Networks, Hierarchical Classification

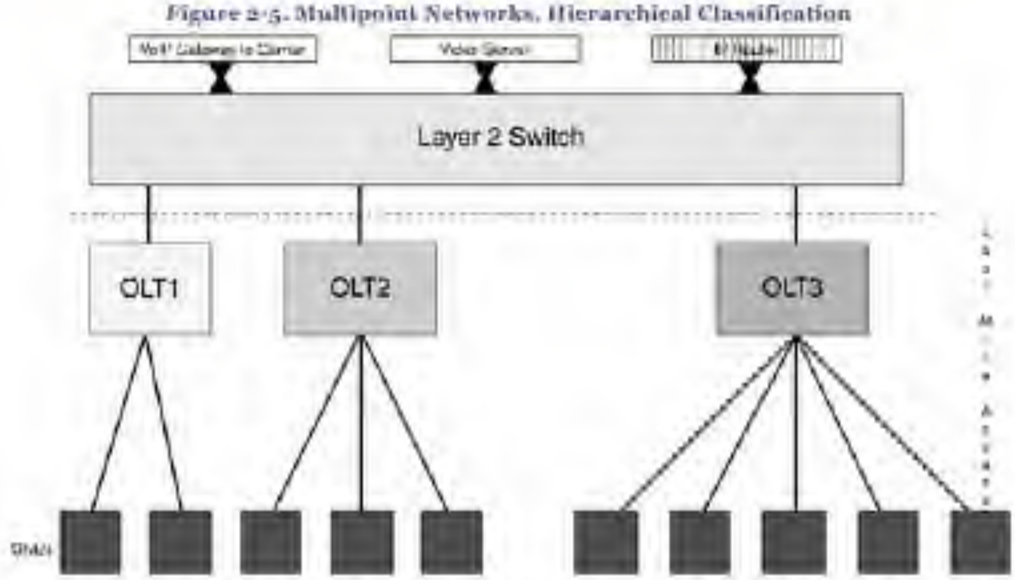

EFM also introduced the concept of EPONs, in which a P2MP network topology is implemented with passive optical splitters, along with optical fiber physical media dependent (PMD) layers that support this topology. In addition, a mechanism for network OAM is included to facilitate network operation and troubleshooting.

EPON is based upon a mechanism named *Multi-Point Control Protocol* (MPCP), defined as a function within the MAC control sublayer. MPCP uses messages, state machines, and timers to control access to a P2MP topology. Each ONU in the P2MP topology contains an instance of the MPCP protocol, which communicates with an instance of MPCP in the OLT.

At the foundation of the EPON/MPCP protocol lies the P2P emulation sublayer, which makes an underlying P2MP network appear as a collection of P2P links to the higher-protocol layers (at and above the MAC client). It achieves this by prepending a logical-link identification (LLID) to the beginning of each packet, replacing two octets of the preamble.

In this book, due to the nonstandardization of the MPCP protocol, we cover two very similar protocols: TUR and IPACT.

Transmission upon Reception

The key to successful communication in access PONs is to provide a solution that alleviates issues of throughput and fairness for upstream communication. The approach has to be low cost and easy to deploy. To alleviate issues faced by TDM and WDM in PON, we propose a protocol called Transmission Upon Reception (TUR). This protocol ensures both downstream and upstream communication and is able to accommodate the variations in traffic between nodes and guarantee fairness to the end-user nodes. TUR also ensures a collision-free system. TUR is a simple protocol that guarantees each end-user node some time to transmit data; the amount of time it gets is proportional to the amount of traffic it gets (from the central office [CO]), with some variations for ensuring fairness.

The idea is as follows: An end user, upon receiving a frame at time t_1 of size S_1, gets the transmitting time from time $t_1 + S_1$ to time $t_1 + S_1 + S_1 - d$, where d is some constant of the system. The constant d neutralizes the differences in propagation delays among different nodes in the system (analogous to the far-near problem). The protocol thus is simple and intuitive but can create undue fairness issues for nodes that have frames to transmit but none to receive. As can easily be observed, there is a "lag" between downstream and upstream communication. (You can see details of such in Figure 2-6.) Consider, for example, a 10-node system: If a node N_x receives a burst of duration T_x at system initialization, it would be able to transmit for duration $T_x - d$ starting from time T_x.

Figure 2-6. System Block Diagram for PON

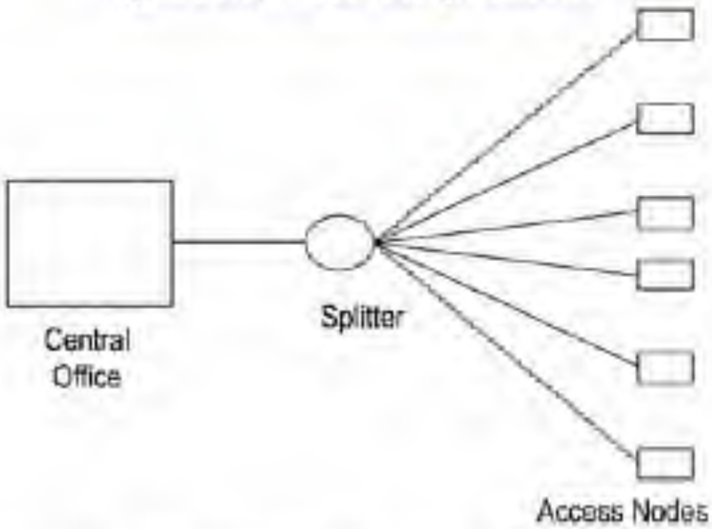

If the next burst of duration T_y is such that $T_{y/} < T_x$, and is destined for node N_y, node N_y can begin transmission from time $T_x + T_y$ to time $T_x + T_y + T_y - d$. However, node T_x is transmitting data from T_x till $2T_x - d$. If $2T_x - d > T_x + T_y$, there would be a collision of the two frames, although we are still adhering to the basic principle of TUR. To avoid this issue, we propose the following corollary to TUR: At initialization, the CO sends a counter called the *start point indicator* to each end user, whose value is zero. As downstream transmission begins, each node starts incrementing the start point indicator by the length of the latest received frame, whether or not this frame is destined to itself. This way every end user knows the status of the upstream channel in terms of the time when there is no transmission. If node N_x receives a burst of frames of duration T_x, all the nodes (inclusive of N_x) increment the counter value to $T_x + T_x$. This means that the earliest value at which the upstream channel is free is $T_x + T_x$. Subsequently if node N_y receives the next burst of frames of duration T_y, all the nodes increment the counter value to $2T_x + T_y$. Also, because the frames were destined for node N_y, the node gets the right of transmission from $2T_x$ to $2T_x + T_y$, although it ends its transmission at time $2T_x + T_y - d$ for reasons explained later. TUR basically allots time slots (of variable length) to nodes on the upstream channel. This corollary to the basic TUR protocol ensures collision-free access of the channel but creates some voids in the upstream channel. The voids are created when the present burst of frames is greater in duration than the previous burst of frames (to a particular user) and this difference exceeds the lag time between upstream and downstream channels. Because we assumed a burst protocol as opposed to a frame-based protocol, the bursts are quite large compared to the voids, and from simulation we observed that voids represented just 3 percent of the channel capacity even after transmission of 10 million bursts (of Ethernet frames). Note here that TUR is not a TDM protocol in the real sense but is a statistical

First Mile Access Networks and Enabling Technologies By Adrian Gonzalez, Tony Antony
ISBN: Publisher: Cisco Press
Print Publication Date: 2004/02/24

TDM scheme, which does not require synchronization of the end-user nodes with the CO or each other.

Concerning multicasting, if a burst of frames is destined for two or more end users in the downstream direction, then, as per TUR, all the receiving nodes will attempt to transmit from the earliest time the upstream channel is free. To avoid collision due to multicasting, the CO sends a multicasting frame and scheduling information for each member of the multicast group at the beginning of the burst. The end users that are part of the multicast group transmit at their turn. For applications such as video on demand, the members of the multicast group may not even be required to send any information (except small acknowledgements) in the upstream, and hence upstream capacity during this interval can be devoted to some other node (procedures for such are discussed later).

Figure 2-7 shows a receiver block diagram. Note PD has a built-in 1310 filter with a band separator.

Figure 2-7. Receiver Block Diagram

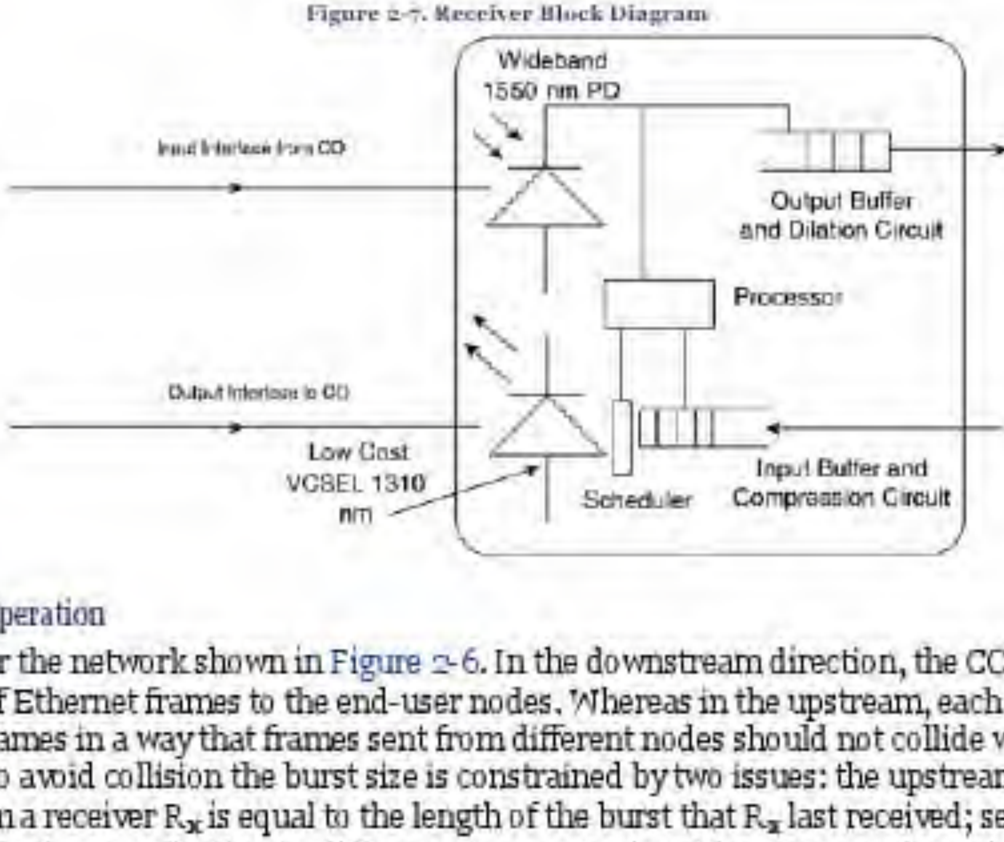

System Operation

Consider the network shown in Figure 2-6. In the downstream direction, the CO sends bursts of Ethernet frames to the end-user nodes. Whereas in the upstream, each end user sends frames in a way that frames sent from different nodes should not collide with each other. To avoid collision the burst size is constrained by two issues: the upstream burst size from a receiver R_x is equal to the length of the burst that R_x last received; secondly the size is also curtailed by the difference in propagation delays from each receiver (end user) to the CO. At time t_1 if end user R_x receives a burst of frames of duration t_b, it can begin transmitting from $t_{stpt} + t_b$ and continue till $t_{stpt} + 2t_b$, where t_{stpt} is the start point indicator mentioned in the preceding section. To avoid contention, however, we have to

note that the burst would reach other nodes at different times depending on the propagation delays from the CO to the different nodes. Let us assume that the node we are looking at is not the closest node from the CO. Let the difference in propagation delay between the incumbent node and the node closest to the CO be _. Hence node R_y (closest to CO) would get the frame from t_1 – _, and the frame would last till $t_1 + t_b$ – _. Under the tenets of the TUR protocol, node R_y would assume node R_x sends a burst from time t_{stpt} – _ to time t_{stpt} – _ + t_b, although R_x actually sends a burst from t_{stpt} to $t_{stpt} + t_b$. Hence there is the likelihood of collision. To avoid collision, we reduce the upstream transmission time by a factor of 2_, which essentially takes into consideration the difference in propagation delays from the CO to other nodes.

Let t_i be the amount of time node i gets a burst. Let $\_{j}$ be the difference in propagation delay between node i and the CO and node j and the CO, where j is the node closest to the CO. Then the total amount of time node i can transmit is $TX_i = t_i$ – 2_j. If F is the average frame length (= 1500 bytes in the case of Ethernet), $N_i = t_i/F$ is the number of frames that were transmitted to node i from the CO. Then the total transmission of node N_i is as follows:

$$D = Ct_i\left[1 - \frac{2\delta}{F}\right]$$

(C is the line rate.) Because _ is negligible as compared to frame duration:

$$D \cong Ct_i$$

Here even if we assume the closest end user to be 1 km from the CO and the farthest end user to be 20 km away, _ would still be a fraction of the transmission of an Ethernet frame, and hence the protocol is very tight in terms of bounds on the ratio of upstream to downstream communication.

TUR allots upstream channel access to a node that has received a burst of frames. This may seem to result in long queues for nodes that do not receive sufficient frames from the CO. If so, this could lead to a fairness issue for upstream access among the end users. We solve this issue as follows: The CO has an array of counters with a counter for each end user. The counter notes the time since the last transmitted frame to this particular node. The CO, upon realizing that a node has not received any frames for some time t_{idle}, sends a "status" frame of length L_{status} to this node. The node sends back to the CO the state of its output buffer. With this exercise, the CO is periodically able to know the status of buffers and is able to intelligently multiplex frames, both real and dummy, to ensure some degree of fairness. The drawback, of course, is when a node receives frames but has no frame to send and hence wastes upstream bandwidth. We solve this issue by allotting "empty-buffer" frames to the end user nodes. An end user upon getting bursts of frames can inform the CO that its buffer is empty. For the next few frames (less than the QoS requirement),

the CO tags a "do not send" label behind the upstream frame to this node while sending an "out of turn send" label to an end-user node whose output buffer condition is comparatively severe. In this way, the protocol guarantees degrees of fairness and throughput to the network. The protocol can also be provisioned for SLAs, guaranteeing some basic QoS requirements to an end user.

Interleaved Polling with Adaptive Cycle Time

Another academic protocol proposed for creating high-efficiency collision-free communication in PONs is called *Interleaved Polling with Adaptive Cycle Time* (IPACT). The main principle of IPACT is as follows: An OLT is the central point of intelligence. An OLT polls a grant message to an ONU allowing it to transmit data. Simultaneously, the OLT in the downstream transmission sends data to the ONU. At the end of transmission to a particular ONU, the OLT so plans the upstream transmission that this incumbent receiving ONU informs the OLT of the size of data it has in its buffer. This way the OLT knows how much data is there in each buffer of every ONU.

This protocol is superbly designed for centralized management. However, the protocol lacks in being able to cater to the dynamic variation of IP traffic. Note that IPACT also cannot serve QoS-sensitive traffic very well, which in contrast is taken care of by TUR. Second, in the downstream the OLT has to synchronize the data size being sent to an ONU while granting a request. At the end of the transmission, the ONU responds with its buffer size information. This means that the other ONU, which currently had access to the grant from the OLT, is now transmitting data and, hence, a collision could occur between the transmitting data of the granted ONU and the transmitted control frame of the incumbent ONU. Although IPACT is essentially a STDM protocol, it works in these cycles of data transmission interleaved with grants and other control information. The good aspect of such communication as seen even in TUR is the absolute nonrequirement of a control layer, which is submerged within the data transmission.

Numeric Evaluation of EPONs

We conducted a simulation program to demonstrate the operation of the TUR protocol with bursts of Gigabit Ethernet frames. We assumed bursts rather than single-frame transmission because IP traffic is bursty and, second, because the protocol is intended for burst transport rather than frame transport to maximize the performance. Although, of course, Ethernet frames are transported and the performance is not degraded even if we consider just pure Ethernet frames (assuming fast vertical cavity surface emitting laser [VCSEL] technology at the end users).

Figure 2-8 shows the working of the TUR protocol for four end-user nodes.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 1587051249 Publisher: Cisco Press
Print Publication Date: 2004/02/24

Figure 2-8. Working of the TUR Protocol for Four End-User Nodes (A, B, C, and D)

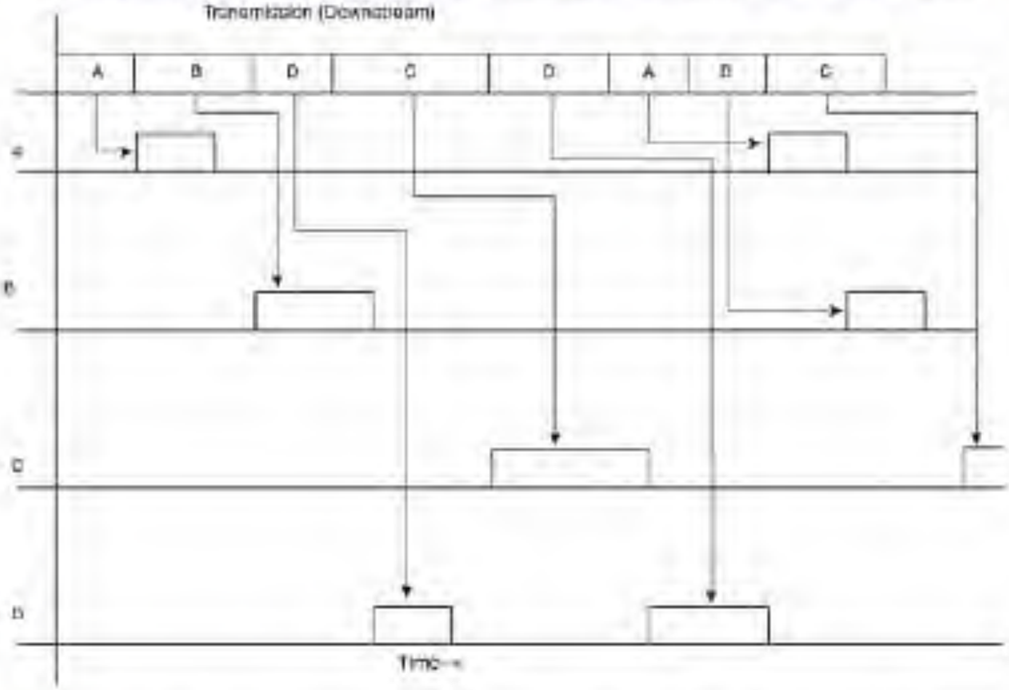

Figure 2-9 shows the efficiency of TUR for different nodes. We assumed a 1-GB Ethernet line rate in both upstream and downstream directions. Frames are generated in a way that bursts of several frames are destined for a randomly selected end user. The frame generation phenomenon for both the CO and the end users follows a Pareto distribution. The mean number of frames that make up a burst in our simulation was 20, and the variance of burst size was 16. We measured throughput of the system. Throughput is defined as the ratio of frames accepted by the end users for transmission to the CO, to the total number of frames that come to the end users. The buffer status of the end users was also measured for various configurations of number of nodes. We observed that the system performs particularly well for 10 to 16 nodes. The buffer requirement was also quite low for that number. As the number of nodes exceeded 16, we observed severe penalty for throughput as well as fairness at the end users. By increasing the line rate proportionally, this issue can be alleviated. The percentage of voids created was less than 3 percent for 10 million frames.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

Figure 2-9. Efficiency of the Protocol with the Number of Nodes and the Corresponding Receiver Buffer Status for a Measurement of H = 0.9

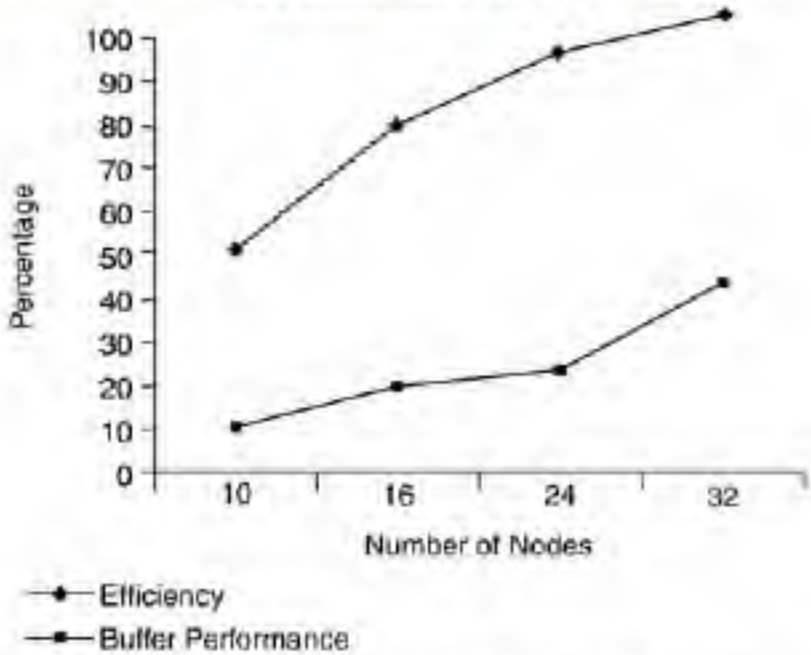

Because we were considering bursty traffic (predominantly IP), we measured the receiver buffer performance for several degrees of burstiness, given by the Hurst parameter in Figure 2-10. As the burstiness of the input source increases, the buffer requirement increases nonlinearly. Note that the buffer requirement does not increase exponentially as noted by an initial study of self-similar traffic (see Kramer, Mukherjee, and Pesavento).

Figure 2-10. Receiver Buffer Condition for Different Values of Burstiness at the Input Interface

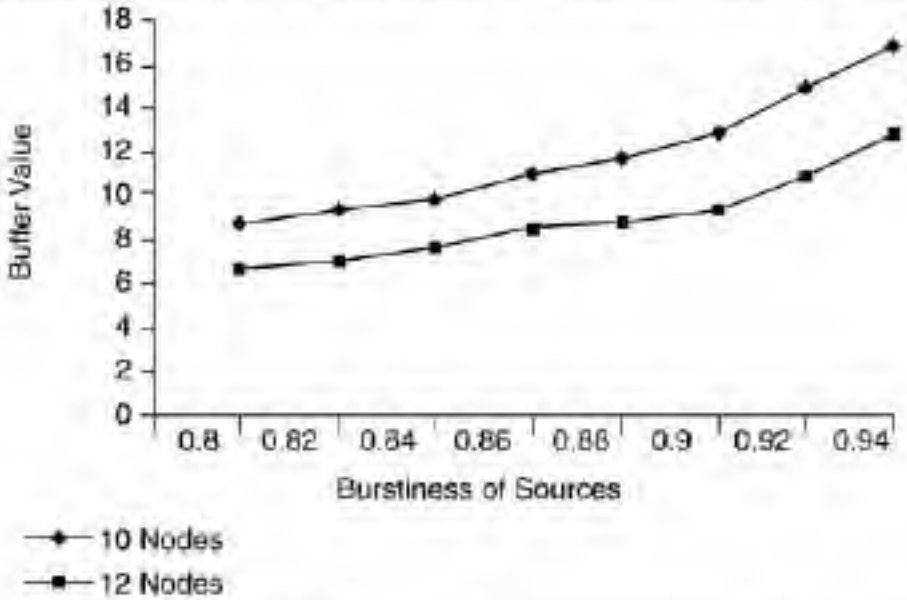

This means that TUR is an efficient protocol for burstycommunication. In Figure 2-11, we have assumed the correction for fairness to be present during simulation. In Figure 2-11 we consider the effect of having correction in TUR for cases when a receiver gets a burst of frames but has no frames to transmit. For normalized loads in the range of 0.3 to 0.75, we observed maximum benefit by using correction in TUR.

Figure 2-11. Blocking Probability in Terms of Dropping of Frames for Schemes With and Without Correction for Normalized Load

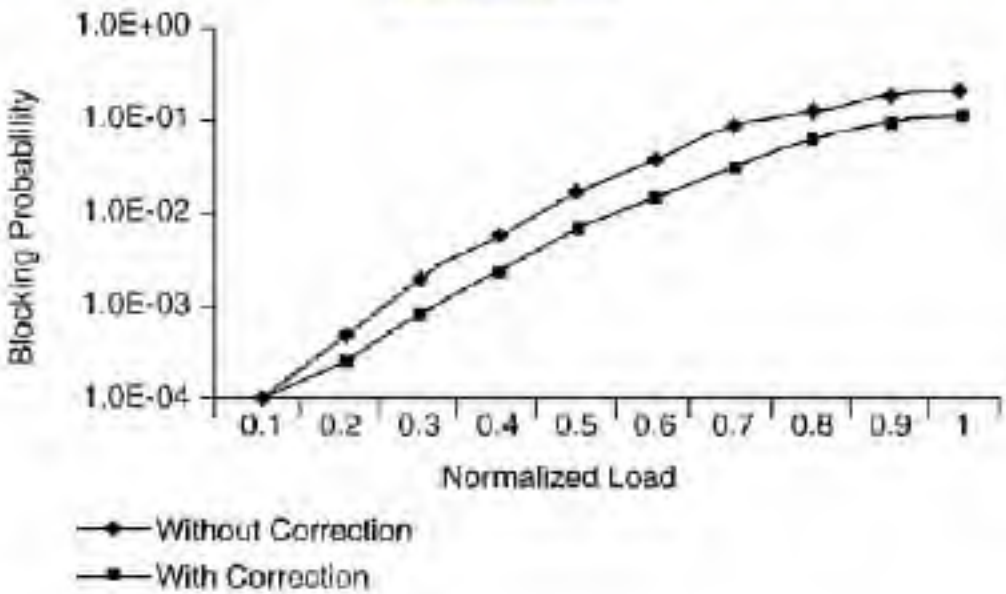

The blocking (percentage of packets dropped at the end user) is reduced by 20 percent in these cases. The number of end users in the measurement of Figure 2-11 was 12, and we observed a similar result for 10 and 16 nodes.

Figure 2-12 shows the fairness issue of TUR in PONs. For different numbers of nodes, we measure fairness for very bursty (H = 0.92) and comparatively less-bursty (H = 0.8) traffic. Fairness here is measured as follows: It is the ratio of times a node gets the upstream bandwidth to the times it has bursts to send scaled and normalized traffic with respect to a TDM scheme. In other words, if for a TDM scheme with 10 nodes a node gets every tenth slot for upstream communication, correspondingly we measure the performance for our scheme with more realism by considering pure bursty traffic.

Figure 2-12. Delay Performance of TUR and IPACT

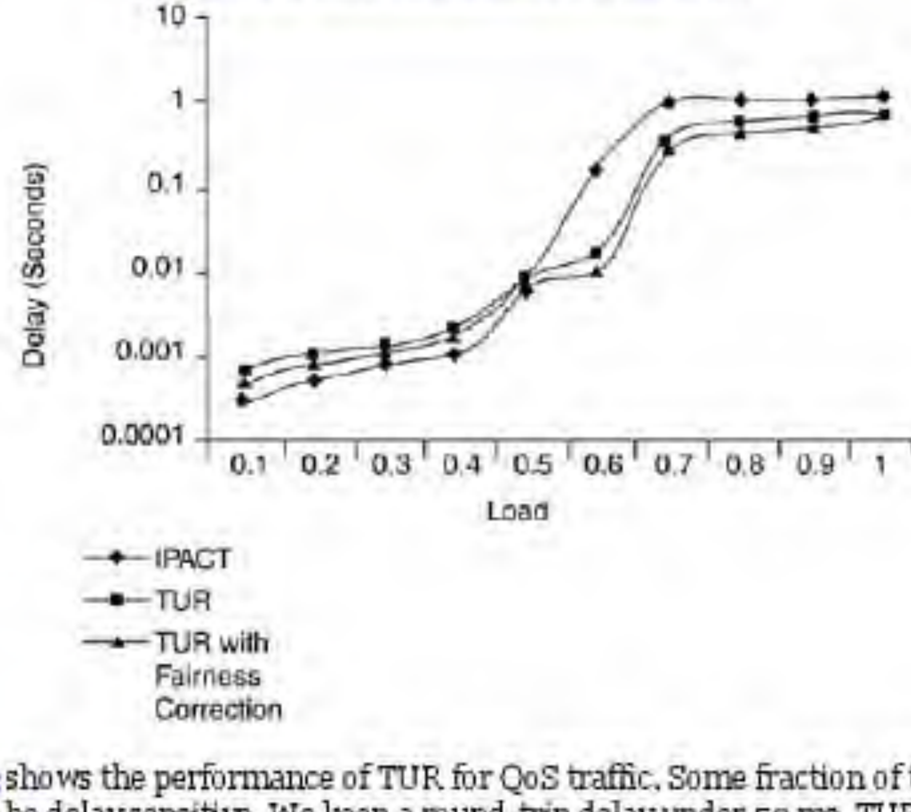

Figure 2-13 shows the performance of TUR for QoS traffic. Some fraction of the traffic is assumed to be delay sensitive. We keep a round-trip delay under 50 ms. TUR performs explicitly well for QoS schemes, too. If we compare TUR versus a TDM protocol for QoS issues, we would expect to find that TUR is initially less proactive to delay requirement, but as loads increase TUR performs better than TDM schemes. This is because through simulation we observed that for a 16-node system, the percentage of empty slots for TDM was about 14 percent. On the other hand, the percentage of empty slots in TUR was merely 2.9 percent. TUR makes good use of the available bandwidth and maximizes network use.

First Mile Access Networks and Enabling Technologies By Adrian Gumaste, Tony Antony
Publisher: Cisco Press
Print Publication Date: 2004/02/24

Figure 2-13. Fairness with Correction for Various Bursty Schemes

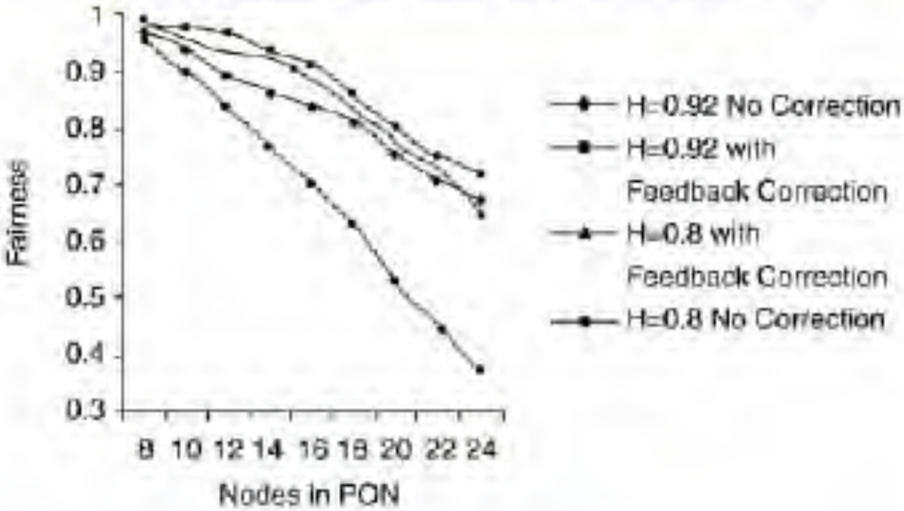

Comparison with Contemporary Schemes

The proposed TUR protocol for access PON is efficient and low cost and is built from very mature technologies. A number of solutions have been proposed that consider various protocols and technologies for implementation. Among them, three generic solutions that prompt interest for comparison are pure WDM, pure TDM, and IPACT.

Table 2-1 shows comparative performance with other PON schemes.

Table 2-1. Comparative Performance with Other PON Schemes

Feature	TUR	IPACT	Pure TDM	Pure WDM
Synchronization for end-user nodes	No	No	Yes	No
Cost	Low	Medium	High	Very high
End-user delay	Medium	Medium	Low	Very low
Throughput	High	Medium	Variable	Underutilized
Scalability	Good	Average	Bounded scalable	Very good (depends on the number of wavelengths)
Efficiency to burstiness in IP traffic	Yes	No	No	No
Enabling technology	Mature	Quite futuristic	Mature	Mature but expensive

In Gumaste and Chlamtac's "A Protocol to Implement IP Centric Communication in EPON," we see a very similar PON implementation with the CO polling different users and creating a pseudo-TDM scheme based on end-user feedback. This scheme, called IPACT, is an efficient and scalable solution. Figure 2-14 compares the delay requirements of frames at the end users using IPACT and TUR schemes respectively. As shown, the delay

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 158705129X Publisher: Cisco Press
Print Publication Date: 2004/02/24

performance of TUR and IPACT is quite similar, but TUR has a lower delay requirement than IPACT for medium to high loads. Load is calculated based on the burstiness of the source and the ratio of generated load to total capacity of the system. Table 2-1 also compares the performance of TUR with pure WDM, pure TDM, and IPACT. To implement dynamic variations in traffic, we see TUR represents a good solution for PON. The technology is very mature, and implementation costs are relatively low. From the table, you can see that the proposed solution is quite in conformity with contemporary protocols.

Figure 2-14. Performance of TUR for QoS Requirements

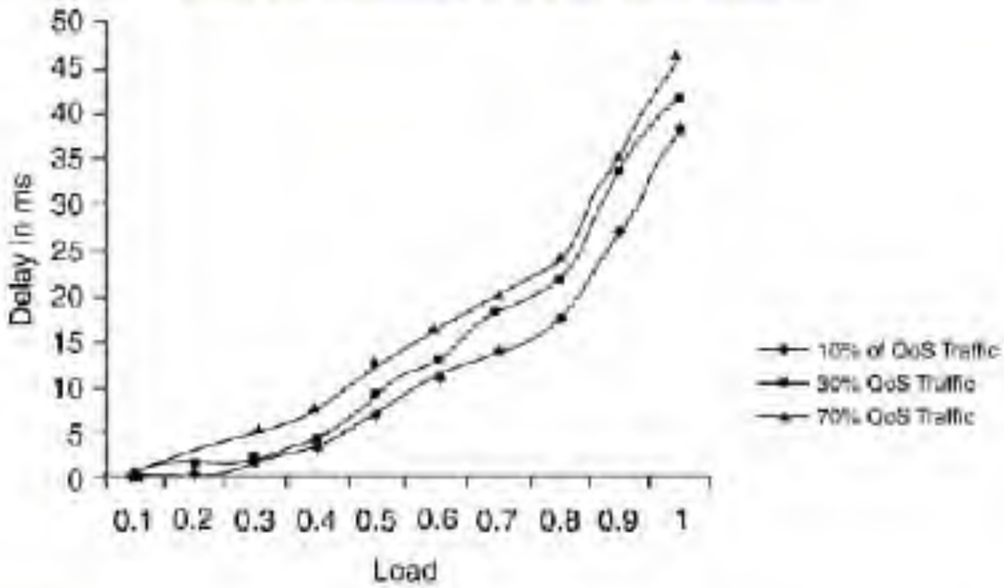

Summary

In this chapter we discussed the multiple PON technologies (including APON, GPON, and EPON) and elements that affect first mile access networks. This chapter also covered the benefits of EPON as compared to APON and GPON. This chapter focused on the aspects of cost, performance, and protocol issues and, therefore, outlined EPON as the key technology of the future in the first mile. In this chapter, we demonstrated the disruptive aspects and protocol behavior of Ethernet. The chapter concluded with an in-depth look at an academically proposed protocol called TUR.

Review Questions

1: Differentiate between the three types of PONs and explain the advantages of each.

2: Explain the advantage of EPON over GPON and BPON.

3: Compare BPON and GPON for a network in which there is a wide differentiation of services.

4: Explain how you would map voice and video in the same physical transmission layer using GFP.

5: What are the key management issues in EPON?

6: How is Ethernet a disruptive technology, and how can you increase its advantage as an incremental service?

7: With regards to cost, compare EPON with GPON and BPON for the same line rates and identical network sizes.

8: Calculate the efficiency of a PON protocol using TDM for bursty traffic. Develop an analytical formula for evaluating the efficiency of the protocol.

9: Based on Question 8, develop a correction protocol that enhances the efficiency of the system by using some sort of adaptation for bursty traffic.

10: Discuss the differences between TDM and STDM and show, with a numeric example, how STDM is more suited for IP-centric traffic.

11: For a 1:32 coupler, if the input power is 0 dB, calculate the power at each ONU that is 10 km away from the coupler. Assume 0.2 db fiber loss.

12: Discuss the far-near problem in PON.

13: Two ONUs are 10 km and 14 km away from an OLT. How much is the differential delay that the further ONU must cater to so that the far-near problem is solved and both ONUs appear equally far from the OLT?

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: 2004/02/24

14: Design a last mile network with 12 end users and 1 central office. The bandwidth requirement for each end user is 3, 7, 9, 11, 1, 2, 4, 16, 32, 10, 8, and 16Mbps respectively. Label these users A to L. Users A, C, F, G are mobile and have to be connected through a wireless LAN. The remaining users are fixed. Optimize the network such that each ONU can cater to 32-Mbps traffic. How many ONUs are needed? From each ONU there is a distribution network to the end users. Optimize the placement of the ONU. Now assume arbitrary distances (min = 8 km, and max = 20 km) for each end user from the OLT. Develop a linear program to maximize performance of the system. Discuss the problems of BPON if used in this system. Now solve those problems faced when using BPON by using TUR in EPON. Compare numeric results.

References

Casson, et al., "Ethernet in the MAN-A Case Study in Disruptive Technology."

Chlamtac, Imrich and Ashwin Gumaste. "Bandwidth Management in Community Networks." Keynote address, International Workshop on Distributed Computing (IWDC). Calcutta, India: December 2003.

Green, Paul, "Fiber to the Home" Keynote address, Opticomm 2002 in association with ITCOM, Boston: 2002.

Gumaste, Ashwin and Imrich Chlamtac, "A Protocol to Implement IP Centric Communication in EPON," Proc of Optical Network Design and Modelling, ONDM, Budapest, Hungary: 2003.

Gumaste, Ashwin and Imrich Chlamtac, "A Protocol for EPON Implementation," Proc of ICC, Anchorage, Alaska: 2003.

Kramer, Glen, et al., "IPACT: A Dynamic Protocol for an Ethernet PON (EPON)," *IEEE Communications Magazine*, Volume 40: No. 2: February 2002.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: [illegible] Publisher: Cisco Press
Print Publication Date: [illegible]

Pesavento, Gary, "Ethernet Passive Optical Networks," *IEEE Communications Magazine*, Volume 40: No 2: February 2002.

Tolley, Bob, "Point-to-Point GBE and OAM for PON," Networld + Interop, Chicago: 2002.

First Mile Access Networks and Enabling Technologies By Ashwin Gumaste, Tony Antony
ISBN: 1587051257 Publisher: Cisco Press
Print Publication Date: 2004/02/24